

DRAFT

DO NOT CITE OR DISTRIBUTE WITHOUT PERMISSION

THE STUFF OF LIFE

TOWARDS A PHILOSOPHY OF MATERIAL CULTURE

© 2006 Beth Preston

Department of Philosophy
University of Georgia
Athens, GA, USA 30602-1627
706-542-2819 (Phone)
706-542-2839 (FAX)
epreston@uga.edu

CHAPTER ONE

THE CENTRALIZED CONTROL MODEL OF PRODUCTION

I was central.
I had control.
I lost my head.

R.E.M., "Country Feedback," *Out of Time*

INTRODUCTION

Perhaps the most obvious fact about material culture is that it is what it is and where it is largely as a result of human activity. So understanding material culture is inextricably bound up with understanding the activities involved in its use and production. In contemporary Anglo-American philosophy there is a specialized area of research called action theory that traces its descent back to Aristotle. Action theorists have not directed their attention specifically to material culture, but an obvious move would be simply to adopt the conceptual framework of current action theory as the foundation for a theory of material culture. Randall Dipert (1993) did just that in his path breaking action-theoretic account of artifacts. The burden of this chapter is to explain why, contrary to apparently reasonable expectations, this is actually a very bad move.

First, action theory has focused almost exclusively on the actions of individual agents. It is only in the last ten years or so that a small body of literature on joint intentions and actions has grown up.¹ But items of material culture are not only typically produced collaboratively; they are typically used collaboratively. No one makes a car by themselves; and while driving we are constrained to collaborate with other car users on public roads, as well as sometimes with other passengers in our own vehicle. Even something as routine as making dinner is very often a collaborative enterprise, with regard to both production and consumption. So the first problem with action theory as a foundation for a theory of material culture is the lack of a well developed account of collaborative action.

Second, action theory has concentrated exclusively on the analysis of intentional action, and in recent years it has swung rather heavily towards planning as the fundamental concept in its analysis of intention. Plans do have an important role to play in human action, but not all our activity is planned in the relevant sense. Rather much of our everyday activity, including our interactions with items of material culture, has an improvisational structure. For example, you cannot plan ahead for every action you will perform while driving from one place to another because you cannot know what other cars and pedestrians you will encounter, not to mention

¹ This philosophical literature was preceded by a body of literature on multiple agent planning in artificial intelligence. So work on joint intention and action might be considered to have begun twenty years ago rather than only ten. But the connection between these bodies of literature is rather tenuous, so it does not seem appropriate to consider them as a single body of work in spite of some similarity in the subject matter.

roadwork, downpours, or deer. Similarly, the production of dinner is often more of an improvised than a planned activity. For example, you may come home with no particular plans for dinner, look in the refrigerator, and put together a meal from basic ingredients and leftovers. Or you may look in the refrigerator, not find anything suitable or appealing, and decide to go out for dinner. So the second problem with action theory as a foundation for a theory of material culture is the lack of an account of improvisation to complement its already well developed account of planning.

A corollary to this second problem is that items of material culture are not always the direct result of relevant intentions. A simple example is the unintended production of a path by people intent only on getting from one place to another. More complex examples have to do with the fact that material culture is an active and constitutive aspect of social institutions, and social institutions are typically underdetermined by the intentions of either their creators or their users. Foucault's (1977) analysis of Bentham's model prison, the Panopticon, is an excellent example of how material culture has significance that informs our activities over and above any intentions or even self-understanding we may have. So it is not just the growing commitment to a planning theory of intention in recent action theory, but the prior restriction of the focus to intentional action that is problematic.

Action theory thus lacks some crucial resources as a foundation for a theory of material culture. But the features that generate this problem are not entirely a matter of recent innovation. They have roots in the history of action theory, starting with Aristotle. More specifically, over the course of its long history action theory has consistently focused on and developed a model of action I shall call the *centralized control* model. The dominant features of recent action theory outlined above represent the most recent and most refined version of this model. The problems posed by action theory for a theory of material culture are a consequence of the one-sidedness of this model, which overemphasizes some aspects of action while marginalizing other aspects or leaving them out of account altogether. Of course, it might be argued that this marginalization is principled; that the aspects of action that have been ignored are secondary phenomena, or relatively unimportant for some other reason. But in point of fact no such arguments are ever made. Moreover, it is part of my project to show that the marginalized and ignored aspects—improvisation and collaborative activity in particular—are foundational phenomena of human action the investigation of which cannot be postponed or subordinated without introducing serious distortions into the theory.

I shall begin this chapter with a brief history of what might be called production theory. Aristotle, Karl Marx, and Randall Dipert all single out the production of items of material culture as a human activity deserving of study in its own right. And as Marx points out, human production involves the use of existing material culture (tools, workplaces, materials that have been worked up to some extent and so are not completely “raw,” and so on).² Thus the study of production is at the same time the study of use. Most importantly for our purposes, all of these authors understand production in terms of the centralized control model predominating in action theory in general. So examining their accounts of production will serve the dual purpose of explicating the development of the current planning oriented version of the centralized control model and simultaneously providing an historical overview of the (unfortunately very meager)

² See Marx (1976), especially Chapter 7.

extent to which action theory has been applied to the study of material culture. The second half of this chapter will then examine the issues of collaboration and improvisation in more detail in the context of the planning version of the centralized control model.

ARISTOTLE

For Aristotle, production consists in impressing a form on matter. We produce neither the form nor the matter, but only the union of the two (*Metaphysics* 1033a24-1034b5). This requires thinking about how to bring the desired form into union with suitable matter.

Every art [*techne*] is concerned with bringing something into existence, and to think by art is to investigate how to generate something which may or may not exist and of which the [moving] principle is in the producer and not in the thing produced.... (*Nicomachean Ethics* 1140a12-15)

Things generated by art are those whose form is in the soul.... Now the healthy is generated when a man thinks as follows: since health is so-and-so, if the subject is to be healthy it must have such-and-such, let us say uniformity, and if uniformity, then warmth; and he always thinks in this manner until he arrives at something final which he himself can produce. Then the motion from this instant onward, which here is a motion towards health, is called “production” [*poiesis*]. Thus, it turns out that in a sense health is generated from health, and a house from a house (that is, the material house from the house without matter), for the medical art and the building art are the forms, respectively, of health and of the house....

Of the generations and motions just considered, one of them is called “thinking” [*noesis*] and the other “production” [*poiesis*]; thinking occurs from the principle or the form, production from the end of thinking and thereafter. (*Metaphysics* 1032b15-20)

The first thing to notice here is that the process starts with a specification of the thing to be produced (the form in the mind of the producer) and the ensuing deliberation or thinking concerns the specification of the steps by which this form may be realized in appropriate matter. So at the end of the thinking process the producer has in her mind a mental design for the product, complete with step-by-step instructions for constructing it. Second, Aristotle suggests that this mental design is finished prior to the production proper, the actual construction. So for Aristotle there are two clearly demarcated phases in the overall production process—an antecedent design phase and a subsequent construction phase. Moreover, since all of the thinking is relegated to the design phase, the construction phase must be a matter of unintelligent execution. In construction, the design is simply realized by faithfully executing the embedded step-by-step instructions. So for Aristotle the real interest of the overall production process lies in the mental process of design, not in the actual construction of the item of material culture during production proper.

Aristotle also has some important things to say about the thinking that goes into the design phase. Art (*techne*), he says, is a state, or habit (*hexis*) involving (*meta*) true reason; lack of art (*atechnia*) is habit involving false reason (*Nicomachean Ethics* 1140a7-25).³ So production requires thinking, but this thinking relies on the existence of established habits that it does not have to articulate on each occasion. For example, in thinking about building a house, the experienced builder (or the designer who relies on an experienced builder to execute the design) does not have to specify step-by-step instructions for using a hammer or putting up drywall. Using a hammer competently is a motor habit expected even of a not very experienced builder;

³ This makes art analogous to virtuous action, which is a habit of choosing in accordance with a mean defined by true reason (*Nicomachean Ethics* 1106b36-1107a2).

and putting up drywall competently is a habituated technique expected of an experienced builder. Thus Aristotle's view is not simplistically rationalistic. Every deliberately generated design will implicitly or explicitly contain elements that are not themselves deliberately generated by the designer during the design phase or by the executor during the construction phase. These habit elements are invoked as integral wholes, and their internal rationality (as opposed to the external rationality of their place in the design, which may require deliberation) depends on the rationality of the already accomplished habit acquisition process.

It is also important to note something Aristotle does *not* discuss—the typically collaborative nature of production. As the passages quoted above from the *Metaphysics* and the *Nicomachean Ethics* show, Aristotle routinely speaks of the producer, usually a skilled artisan or expert, like a potter or a physician, as a single individual working alone. So his account of production is at least implicitly individualistic. This might be regarded as an expository device—a reconstruction of production in terms of the single individual in order to make the account easier to understand. But if this is an expository device it is not without philosophical consequences. In real life, as opposed to philosophical reconstruction, people typically collaborate in producing material culture, and they do so in a number of importantly different ways. For example, the designer is often not the same individual as the constructor of an item. Similarly, either the construction phase or the design phase or both may involve teams of collaborators working together. And the problem is that Aristotle makes no attempt to show how, if at all, his individualistic account of production and action applies to such collaborative activities.

Finally, the most unusual feature of Aristotle's account of production from a contemporary point of view is his insistence that production and action are distinct.

That which may or may not be can be an object produced as well as an object of action. Now production [*poiesis*] is distinct from action [*praxis*] (and one may be convinced of this from public writings), and so practical dispositions with reason are distinct from productive dispositions with reason: and in view of this the two exclude each other, for no action is a production, and no production is an action. (*Nicomachean Ethics* 1140a1-10)

Thus production is not a subspecies of action, but a distinct species of activity in its own right. Nevertheless, production and action *are* clearly the same in one respect—both are “dispositions with reason,” that is, they involve a process of deliberation concerning what means we may employ to attain our ends with regard to things we can change, as described above in the case of production (cf. *Nicomachean Ethics* 1112b25-1113a5, 1140a10-15). Where production and action differ is with regard to the relationship between the end and this reason driven process of attaining it.

...[F]or the end of production is some other thing [i.e., a product], but in the case of action there is no other end (for a good action is itself the end). (*Nicomachean Ethics* 1140b5-10)

Thus the end of production is something external to and independent of the production process; whereas the end of action is internal to the action. This distinction is also connected to a parallel distinction between motion (*kinesis*) and energetic activities (*energeia*).⁴ Examples of motions are knitting a sweater or going on a diet. The goal you wish to achieve—having a sweater or being thin—is achieved only when and as the process of knitting or dieting terminates. In other words, you do not lose weight for the sake of losing weight, but for the sake of being thin at the end of the process of losing weight. Examples of energetic activities are seeing or flute playing. Here the goal you wish to achieve is in effect already achieved the minute you start the activity. Flute

⁴ See especially *Metaphysics* 1048b17-35.

playing or seeing are thus activities you can engage in for their own sake and not for the sake of some result that will occur only at the end of the activity. Energetic activities may, of course, simultaneously be practical activities with further ends. For example, seeing usually has further ends like finding your keys or not bumping into the furniture; and flute playing might also be for the sake of impressing others or earning a fee. But Aristotle's point here seems to be that activities like seeing or playing the flute *can* be done simply for the sake of doing them; whereas it would make no sense to say that you were losing weight just for the sake of losing it and not for some further end like being thin or lowering your cholesterol.⁵

In proposing these distinctions Aristotle is not just pursuing a descriptive project, but an evaluative one. His purpose is not only to understand what different kinds of activities people engage in, and how those activities are generated, but also to rank activities with regard to their worthiness and their propensity for contributing to the good life. As he explains at the beginning of the *Nicomachean Ethics* (1094a1-1094b10) the principle of this ranking is the extent to which we engage in an activity for its own sake. And by this criterion production is clearly inferior to action, because the process of producing is undertaken only for the sake of the product, which is therefore superior to the process (*Nicomachean Ethics* 1094a5-7). So we cannot achieve human excellence in production, because productive activity only aims at the excellence of the product, not at the excellence of the activity itself. But actions, to the extent that they are undertaken for their own sake, aim at their own excellence and are thus good in the very doing of them. Perhaps Aristotle's most telling pronouncement on the inferiority of production is his remark that productive activity must be missing entirely from the divine life of the gods (*Nicomachean Ethics* 1178b20)—certainly a sharp contrast with the Judaeo-Christian view of production as a prerogative of the divine that humans usurp at their peril. In any case, the important point is that for Aristotle the descriptive distinction between production and action has evaluative force. Indeed, it might well be fair to say that Aristotle insists on the descriptive distinction primarily for evaluative reasons.

Aristotle's distinction between production and action has disappeared without a trace from contemporary Anglo-American action theory. Action theorists assume without discussion that production is simply a subspecies of action. For instance, Myles Brand lists building a bridge as a paradigm case of action right alongside raising your arm and buying a loaf of bread (1970, 3). I think the reason for the disappearance of the distinction is that the basis on which Aristotle proposes it, the status of the end with regard to the process, does not have any valence for contemporary Anglo-American action theorists. Descriptively, their concern is exclusively with the nature of the process by which intentional activity is generated; and as we noted above Aristotle takes this process to be the same in the case of both production and action. Moreover, contemporary action theorists do not in general concern themselves with the evaluative issues about the good life which were so central for Aristotle, and which motivated his thinking about action and production to such a large extent.⁶

⁵ On Aristotle's view there really is only one activity that is perfectly energetic, in the sense that it is chosen for itself alone and not for some other effect to be achieved, and that is theoretical activity, or contemplation. Aristotle claims that contemplation is the highest and most distinctively human activity, and is thus engaged in solely for the sake of engaging in it and not with any practical end in view. See especially *Nicomachean Ethics* 1177a10-1178b34.

⁶ Contemporary action theorists are often concerned with specific moral issues, such as the assignment of moral blame, but this is not the same sort of concern Aristotle had for the overall character of a person's life.

In the Continental tradition of political theory, on the contrary, heroic attempts have been made to counteract any such assimilation of production to action and to reestablish the Aristotelian distinction, complete with its original evaluative force, on a new and firmer basis. Such a project lies at the heart of Jürgen Habermas' longstanding critique of Marx's production paradigm, and his own theory of communicative action.⁷ It also underwrites Hannah Arendt's (1958) influential analysis of the *vita activa*, in which she argues for a rigorous three-fold distinction between labor, work (fabrication), and action. This retention and rehabilitation of Aristotle's distinction between production and action by Continental political theorists reflects their ongoing concern for basically Aristotelian issues concerning the good life, transposed to the context of the contemporary political situation and ethical climate.

This raises the question of what role, if any, Aristotle's production-action distinction should play in a basic theory of material culture. It seems clear that it is a significant part of the history of philosophical thinking about material culture, so it should not simply be passed over in silence. On the other hand, accepting it at face value as part of the basic conceptual framework for the analysis of material culture would be a much bigger mistake. This is because the phenomena of material culture cut across Aristotle's distinction in a number of different ways. First, a theory of material culture must consider not just the creation of items of material culture, but also their use. And clearly the use of material culture occurs in both productions and actions in Aristotle's sense. Second, and more importantly, Aristotle's distinction categorizes activities in a very abstract way that is not aimed at capturing anything about their role in creating material culture. Thus many activities we would not regard as creating material culture fall on the production side of the distinction—losing weight, or curing an illness, for example. On the other hand, some activities we might well regard as creating material culture fall on the action side—music or dance, for example. Last but not least, the evaluative dimension of Aristotle's distinction should give us pause. For Aristotle improvising a song is not just descriptively different from painting a painting—it is a *better* and more praiseworthy kind of activity. It would certainly be tendentious to build such value judgements into the basic conceptual framework of a theory of material culture, even if only implicitly. So Aristotle's production-action distinction must be relegated to a purely historical role for our purposes here.

MARX

For Marx, production is the most fundamental form of human activity; the activity that in the first instance distinguishes us from non-human animals.

Man can be distinguished from the animal by consciousness, religion, or anything else you please. He begins to distinguish himself from the animal the moment he begins to *produce* his means of subsistence, a step required by his physical organization....

This mode of production must not be viewed simply as reproduction of the physical existence of individuals. Rather it is a definite form of their activity, a definite way of expressing their life, a definite *mode of life*. (Marx 1967, 409)

Marx does not mean to deny that other animals produce things; rather it is *how* we produce that distinguishes us from them.

The animal is immediately one with its life activity, not distinct from it.... Man makes his life activity itself [i.e., production] into an object of will and consciousness. He has conscious life activity. It is not a

⁷ See Grumley (1992) for an overview.

determination with which he immediately identifies. Conscious life activity distinguishes man immediately from the life activity of the animal.... Only on that account is his activity free activity. (Marx 1967, 294)

As Marx goes on to explain, this means, first, that while non-human animals produce only to satisfy their physical needs, we produce largely in freedom from basic physical needs; and second, while non-animals produce according to species-specific standards, we produce according to freely variable standards of our own devising, and in particular according to aesthetic standards (Marx 1967, 294-5). Although Marx does not say so explicitly, it seems clear that these standards are for the most part cultural standards, not individual standards. Thus production takes on different forms under historically local conditions, and the history of these changing forms is the history of what Marx calls “modes of production.” In short, on Marx’s view human beings, unlike other animals, do not just produce a few things for their immediate personal use, but produce material cultures that subserve social institutions and are passed down from generation to generation. But the main point for our purposes is that this is all possible because human production is under conscious, rational control and thus free rather than under the control of instinct.

In a famous passage Marx goes on to suggest that this conscious, rational control has a core structure that is species-specific and universal.

We presuppose labour in a form in which it is an exclusively human characteristic. A spider conducts operations which resemble those of the weaver, and a bee would put many a human architect to shame by the construction of its honeycomb cells. But what distinguishes the worst architect from the best of bees is that the architect builds the cell in his mind before he constructs it in wax. At the end of every labour process, a result emerges which had already been conceived by the worker at the beginning, hence already existed ideally. Man not only effects a change of form in the materials of nature; he also realizes [*verwirklicht*] his own purpose in those materials. And this is a purpose he is conscious of, it determines the mode of his activity with the rigidity of a law, and he must subordinate his will to it. (Marx 1976, 283-4)

The account of production here is strikingly similar to Aristotle’s. Marx distinguishes an antecedent, mental design phase from a subsequent construction phase. The design phase includes not only a representation of the thing to be made but also a step-by-step specification of how to make it in the form of a mental rehearsal of the construction process. The construction phase must, then, be an unintelligent execution of the design, which is fully laid out in advance in the mind of the producer. Marx does not specifically describe the procedure by which the design is formed as a matter of deliberating about the means for achieving a given end. But he does say that the end, or purpose, “determines” the activity of the producer, and that this determination is accomplished through consciousness of the end. So we may reasonably suppose that, like Aristotle, what he had in mind is that the conscious purpose of the producer guides and constrains her thinking about the steps necessary to accomplish it.

There is one important departure from Aristotle’s view here, though. Rather than being a realization of an independent, pre-existing form, the product is expressive of the producer’s own needs and purposes. This marks Marx’s account as a modern rather than an ancient conception of the role of production in human life. More specifically, it is an interested, economic conception rather than a disinterested, aesthetic one. For Aristotle, the producer is merely a calculator who takes pre-existing elements—the form and the matter—and figures out how to realize the one in the other. For Marx, when you make something you do not merely change the form of the material you work on; you also endow it with a use-value corresponding to, and expressive of, your own socially and historically conditioned needs and purposes. Marx’s laborer thus creates

material culture in a way Aristotle's artisan does not, because on Marx's view it is production in the first instance that is responsive to changing conditions, needs, and desires in a society and is thus the driving force behind cultural innovation and progress. All other changes in social institutions and ways of thinking are dependent on these changes in the mode of production.

Conceiving, thinking, and the intellectual relationships of men appear here as the direct result of their material behavior. The same applies to intellectual production as manifested in a people's language of politics, law, morality, religion, metaphysics, etc. Men are the producers of their conceptions, ideas, etc., but these are real, active men, as they are conditioned by a definite development of their productive forces and of the relationships corresponding to these up to their highest forms. (Marx 1967, 414)

Moreover, as this passage also shows, Marx does not recognize Aristotle's distinction between production and action. Instead, he characterizes all human activity as production, but sees the production of material culture as fundamental, and all other kinds of production as peripheral and dependent. This also has the effect of reversing Aristotle's evaluation of production and action. Not only is production for Marx the most significant and valued human activity—it has effectively swallowed up all other kinds of activity, which are now just secondary forms of itself.⁸

Marx, like Aristotle, routinely speaks of the producer as a single individual. But in Marx's case there is some reason for thinking this is more than an expository device. There is a difference between conceiving of production as inherently social—which Marx does—and conceiving of it as inherently collaborative, about which he is ambivalent, at best. For Marx, production is inherently social, first, because the individuals who engage in it are the individuals they are only in virtue of the social conditions in which they have lived and which have informed their development; and second, because the mode of production in which individuals engage is specific to historically local social conditions. But you are fully social in this sense even when *not* collaborating in the literal sense of working together face-to-face.

To be avoided above all is establishing “society” once again as an abstraction over against the individual. The individual *is* the *social being*. The expression of his life—even if it does not appear immediately in the form of a *communal* expression carried out together with others—is therefore an expression and assertion of *social life*. (Marx 1967, 306)

In his early work Marx explicitly valorizes individual production as the ideal form of production, and thus implicitly devalues collaborative production at the same time. This comes out in his brief discussion of what he calls “free human production,” the form production would take in the ideal communist society.

Suppose we had produced things as human beings: in his production each of us would have *twice affirmed* himself and the other. (1) In my *production* I would have objectified my *individuality* and its *particularity*, and in the course of the activity I would have enjoyed an *individual life*; in viewing the object I would have experienced the individual joy of knowing my personality as an *objective, sensuously perceptible, and indubitable* power. (2) In your satisfaction and your use of my product I would have had the *direct* and conscious satisfaction that my work satisfied a *human* need, that it objectified *human* nature, and that it created an object appropriate to the need of another *human* being. (3) I would have been the *mediator* between you and the species and you would have experienced me as a redintegration of your own nature and a necessary part of your self; I would have been affirmed in your thought as well as your love. (4) In my individual life I would have directly created your life; in my individual activity I would have immediately *confirmed* and *realized* my true *human* and *social* nature. (Marx 1967, 281)

⁸ Hannah Arendt (1958) explains (and complains about) this reversal in great detail. See especially Chapter III. On Arendt's view, Marx did not engineer the reversal, which is typical of the political and social theory of the last few centuries in general. But he did express it in a particularly well focused form, and his influence in securing its establishment is undeniable.

This passage glorifies the social aspects of production and the relationship between the producer and the consumer. But if a product is to be the expression of an individual's particularity, then the production process, or at least the design phase, must be an individual activity rather than a collaborative one. To the extent that I collaborate with you in the design of a product, it will not be an expression of *my* individuality, but at best a compromise which does not express the individuality of either one of us. Thus the point at which others enter the picture here is necessarily as consumers of my product, not as co-producers. And although the individual producer both exemplifies social characteristics and uses his product to promote social bonding, his productive activity itself is presented as a solitary activity.

Marx's later views are more complex. He came to realize the importance of collaboration—particularly as it involves division of labor in industrial manufacturing contexts—in increasing productivity and thus underwriting economic and technological progress.

When the worker co-operates in a planned way with others, he strips off the fetters of his individuality, and develops the capabilities of his species. (Marx 1976, 447)

But at the same time, Marx criticizes all historical forms of cooperation for restricting the freedom of action of the individual. There are passages in both earlier and later works which point to the unavoidable division of labor in cooperative endeavors as the underlying cause of this curtailment of individual freedom.

...[T]he division of labor offers us the first example for the fact that man's own act becomes an alien power opposed to him and enslaving him instead of being controlled by him—as long as man remains in natural society, as long as a split exists between the particular and the common interest, and as long as the activity is not voluntarily but naturally divided. For as soon as labor is distributed, each person has a particular, exclusive area of activity which is imposed on him and from which he cannot escape. (Marx 1967, 424)

This applies to collaboration at any scale—from small, face-to-face task groups to large capitalist manufacturing plants to work specialization in society as a whole. The point is that in all historically attested collaborative situations the individual is subject to an external power directing her activity; and this is not necessarily the opposing power of another individual, but the institutional power inhering in the distinctive social structures of the collaborative interaction. This situation is aggravated in the division of labor in large scale manufacturing under capitalism, which dehumanizes the worker by relegating her to the machine-like repetition of unskilled operations.⁹ Marx calls this contemporary form of cooperation where workers collaborate only under the control of the capitalist, “purely despotic” (Marx 1976, 450).

Moreover, the co-operation of wage-labourers is entirely brought about by the capital that employs them. Their unification into one single productive body, and the establishment of a connection between their individual functions, lies outside their competence. These things are not their own act, but the act of the capital that brings them together and maintains them in that situation. Hence the interconnection between their various labours confronts them, in the realm of ideas, as a plan drawn up by the capitalist, and, in practice, as his authority, as the powerful will of a being outside them, who subjects their activity to his purpose. (Marx 1976, 449-50)

Thus far it seems that loss of individual freedom of action is an ineluctable side effect of collaboration in production, especially in industrial manufacturing.

But Marx has a vision. He believes that in the ideal communist society the advantages of cooperation will be maintained, but the freedom of the individual will not be curtailed by “natural” collaborative structures. Rather it will be preserved in a “free association” of workers.

⁹ Cf. *Capital*, Volume I, Chapters 13 and 14, *passim*.

In communist society, however, where nobody has an exclusive area of activity and each can train himself in any branch he wishes, society regulates the general production, making it possible for me to do one thing today and another tomorrow, to hunt in the morning, fish in the afternoon, breed cattle in the evening, criticize after dinner, just as I like, without ever becoming a hunter, a fisherman, a herdsman, or a critic. (Marx 1967, 424-25)

In place of the old bourgeois society, with its classes and class antagonism, we shall have an association, in which the free development of each is the condition for the free development of all. (Marx and Engels, 1999, 85)

Marx was notoriously vague about how this free association of producers would work, and equally vague about how the goods and services produced would be distributed to those who needed them. But he was very clear that the possibility of such a free association depended on the abolition of private property and the institution of collective ownership of all significant social and economic resources. As he explains in Chapter 32 of *Capital*, in the bad old days individual producers owned their own means of production and they were free producers in that respect. But the fragmentation and isolation of resources attendant upon this dispersed individual ownership stifled cooperation, and thus productivity and technological development. At the next major stage, capitalism concentrated ownership of the means of production in a few hands, and turned the vast majority of individual producers into “wage slaves.” But at the same time, capitalist manufacturing systems pooled resources and organized cooperation on a vast scale, thus promoting higher productivity and technological progress. Marx predicts that the next stage, the transition to communism, will return ownership of the means of production to the producers, but this time as collective rather than individual property. Thus individual workers will both own the means of production *and* have the organizational knowledge for large scale cooperation gained over the course of the industrial revolution (Marx 1976, 927-30). Ultimately, when the transition to communism is complete, the whole idea of ownership will drop out, and so will no longer govern our relationships to our material culture, to nature, or to each other. In particular, the exploitation of others through the expropriation of the means of production will become literally inconceivable. This will render individual producers truly free for the first time in history, since for the first time they will be able to collaborate with each other unconstrained by imposed social roles or limitations on their access to the means of production.¹⁰

What Marx is envisioning here is a sort of Hegelian synthesis combining maximum individual freedom of activity (from the bad old feudal days) with maximum social cooperation (from bad old contemporary capitalism). In the communist future, the cooperating producer will nevertheless be able to produce in complete independence, i.e., *as if* working alone. Thus individual production as an *ideal* is still upheld here by Marx—indeed, it is *aufgehoben* in this synthesis, since it is not only preserved but enhanced by its combination with cooperation. So the twist here is that in the ideal communist society Marx thinks it will be possible to have the individualistic cake and eat it too. The individual will be completely free to produce as and when she wishes, and yet society will still reap the benefits of cooperation thanks to some vaguely specified social arrangements for coordinating productive activity and distributing products. Thus even in his later work, after having recognized the importance of collaboration in human production, Marx still holds on to individual production and the freedom of activity inherent in it

¹⁰ Marx articulates this vision in several versions, particularly in the chapter entitled “Private Property and Communism” in *Economic and Philosophic Manuscripts* and in *Critique of the Gotha Program*. The issue of ownership and its disappearance under full-fledged communism is most clearly discussed in “Private Property and Communism.” These works are available from www.marxists.org.

as the ideal to be achieved under true communism, and so strengthens the nascent individualism already evident in Aristotle.¹¹ But unlike Aristotle—and indeed unlike most action theorists in philosophy, ancient or modern—Marx thinks very seriously about collaboration in production in general, and has important things to say about its organization in large scale manufacturing settings in particular. Where he falls short is in his understanding of the relationship between individual productive activity and collaborative productive activity.¹²

DIPERT

Contemporary action theory, which inherits its fundamental approach directly from Aristotle rather than through Marx, has not taken up production as a subject matter in its own right. But Randall Dipert (1993), whose main interest lies in what he calls “artifact theory,” has applied contemporary action theory to the study of artifacts, thus generating an action-theoretic account of production. Dipert's leading idea, which is entirely in keeping with action theory's focus on intentional action, is that artifacts must be understood in terms of the intentions of their creators.¹³

A correct description of an artifact as an artifact describes the artifact in the way that its creator conceived of it—at least as much as is now possible. Specifically, an object is contemplated by an agent, and some of its properties are intentionally modified (or perhaps, intentionally left alone); the production of an artifact is the goal of some intentional activity. (Dipert 1993, 15-16)

Dipert goes on to distinguish between instruments, tools, and artifacts. Instruments are naturally occurring objects that have been intentionally used for a purpose, e.g., a stone used as a hammer. A tool is an instrument that has been intentionally modified for a purpose, e.g., a flint nodule that has been flaked to a sharp edge for use as a hand axe. An artifact is a tool that is intended to be recognized as a tool with a specific purpose. This category, somewhat counterintuitively, includes objects we would normally call tools, such as carpenter's hammers and axes, as well as books, works of art, and so on. It also includes performances, e.g., of musical works, which, Dipert argues, are artifacts fair and square on his definition, even if not commonly thought of as such. Dipert's account of production is meant to apply only to artifacts in this technical sense, although some fairly obvious modifications would easily generate appropriate, parallel accounts of the production of tools and the adoption of instruments.

On Dipert's view, the intentions of the creator of an artifact are the outcome of a “deliberative history” in the course of which the creator contemplates the overall function of the artifact as her end and the possible alternative means for achieving this end, and then forms a complex set of intentions in the form of a construction plan.

Of the many elements of the deliberative history, perhaps most important is the means-ends hierarchy, or plan, according to which the artifactual features were imposed on the object. Elsewhere, I have distinguished among the high-, middle-, and low-level intentions within the plan for an artifact. Roughly,

¹¹ Jon Elster calls this valorization of individual freedom of action and expression “ethical individualism” and says that Marx maintained a lifelong commitment to it. See Elster (1985), especially sections 1.1 and 2.2.7.

¹² Unless, of course, you are willing to buy the Hegelian synthesis he suggests. I would confess to extreme skepticism on this score myself. Interestingly, though, this does seem to be the way cooperative work is organized in ant colonies. There is no central control, and what to do at any particular moment is up to the individual ant. Nevertheless, the cooperative work necessary for the upkeep and reproduction of the ant colony does get done, and done efficiently for the most part (Gordon 1999).

¹³ Dipert routinely uses the term ‘creator’ rather than ‘producer’ or ‘maker.’ The quasi-theological overtones of this choice are not to my taste, but I will follow his usage in this section for the sake of consistency and clarity.

these reflect goals or intentions increasingly subordinated, as means to ends, to the purpose or purposes of the artifact. (Dipert 1993, 54)

My intentions, in making a screwdriver, to have another person come to believe it is a tool for turning screws, or actually come to use it easily for turning screws, are high-level intentions. My intention that the handle be out of stable plastic and that it have a certain shape, that the blade be metal, and so on, are middle-level intentions. They are conceived as means to my high-level intentions as ends. Finally, my intentions that I must use a lathe to create this shape, a drill press to place the hole in the handle for the blade, and so on, are low-level intentions. (Dipert 1993, 151)

Like Aristotle, Dipert emphasizes that this plan, which consists of fully conscious, explicit intentions, is supported in execution by a network of not fully conscious “half-intentional” habits (1993, 49-51). For example, my choice of wood as opposed to plastic for the handle of the screwdriver would be a fully conscious intention, but my manipulation of the lathe in shaping it would be a half-intentional habit.

Dipert’s account of production reproduces the main features we found in both Aristotle’s account and Marx’s account. There is a design phase in which a mental plan for producing the artifact is devised, and then a construction phase in which this plan is carried out. The design phase involves an elaborate hierarchy of intentions, the lower levels of which spell out the steps and conditions necessary for achieving the desired results at the higher levels. This hierarchy of intentions is the result of deliberation about the means for achieving the desired overall end and any intermediate ends which may be required. The most important advance Dipert makes here is to characterize this mental design as a plan. This reflects the influence of action theory, which in the last decade or so has developed a planning theory of intention.¹⁴ Aristotle’s conception of production as involving a mental means-end analysis is already well on the way to a plan based account of production, but the direct appeal to planning theory makes available a much more precise and sophisticated conceptual framework, as we shall see in the next chapter.

Finally, since on Dipert’s view the features of the artifact are imposed on it in accordance with a mental plan, construction would seem to be unintelligent execution. But Dipert introduces an apparently countervailing consideration with regard to plan execution in the case of performance artifacts. The performer is, in effect, the constructor of the artifact. But she is not necessarily its creator, i.e., the original author of the work; and this is particularly true in the case of Western classical music, the domain from which Dipert draws most of his examples. On Dipert’s theory of performance, the performer is an agent whose intention is to carry out the intentions of another agent, the creator (1993, 206). But typically the intentions of the creator of the work are not fully known, and/or there are questions about exactly how to carry them out. So in practice the performer is forced into a creative role, fleshing out the author’s known intentions with some of her own. Thus a performance is almost always to one extent or other a matter of *intelligent* execution, because the performer must make substantive decisions about the artifactual features of the work.

¹⁴ See, for example, Goldman (1970), Brand (1986), Bratman (1987/1999) and Mele (1992). There is also a planning theory of action in artificial intelligence going back at least to Miller, Galanter, and Pribram’s *Plans and the Structure of Behavior* (1960). But although this artificial intelligence literature has probably influenced philosophical action theory indirectly, it is almost never cited by action theorists.

There is some serious doubt whether this theory of performance is really adequate to the phenomena.¹⁵ But for our purposes the important point is that the theory as it stands demonstrates that Dipert *does* uphold the idea of unintelligent execution as an ideal, although one that can never be fully achieved in practice. On his view, if the creator's intentions were completely known the performer would—and should—be an unintelligent executor. So the intelligent execution Dipert allows for is merely *faute de mieux*. Unintelligent execution is what would be expected in the ideal case, and what is aimed for in actual cases. Thus Dipert clings to unintelligent execution as an ideal even while recognizing the importance of intelligent execution in practice, just as Marx clings to individual production as an ideal even while recognizing the importance of collaboration in practice.

Unlike either Aristotle or Marx, Dipert is explicit about the individualism of his view.

What I call "art works" are necessarily experienced as artifacts, and what I call "artifacts" are necessarily conceived in terms of an individual agent, deliberative history, act of creation, and so on. This account of artifacts, and thus of art works, is highly biased toward "individualistic" accounts of agency, thought, and action. I think it is useful to have such an account of idealized individual agency, even if one ultimately ends up rejecting that there are many such examples in actual human behavior. I happen to think the idea is also necessary for making sense of many of our institutions and thoughts about ourselves and that the theory is, to some extent, of some people, and on some occasions, true. (Dipert 1993, 194-5)

Dipert goes so far as to suggest that the deliberative history of all artifacts, even those known to have been created collaboratively, should be reconstructed as if they were the work of a single individual. He mentions that one motivation for this is to give an account of the unity of the artifact in terms of the unity of the producer's intentions as directed towards an end.¹⁶ Thus for Dipert individual production functions as a sort of ideal type in accordance with which actual production activities and their results can be understood. But this leaves his account open to the same worries Aristotle's account raised, since many of the important phenomena of collaborative production may not be captured by such individualistic reconstructions. So however useful these reconstructions may be for some limited purposes, they direct attention away from the study of collaborative production. Thus while acknowledging that individual production is an ideal not always instantiated in practice, and perhaps not even typical of practice, Dipert does not make any significant effort to understand collaborative production or its relationship to individual production, as Marx at least tried to do. On the other hand, Dipert's forthright individualism with regard to production is an important confirmation of the historical tendency towards individualism we detected in less clear forms in both Aristotle and Marx.

THE CENTRALIZED CONTROL MODEL

Each of the three accounts of production we have considered has some unique features. But the basic account of how productive activity is generated and structured is the same in all three cases. All three accounts share the following features:

¹⁵ The main problem is that Dipert ignores the expectation of improvisation in performance that is virtually universal in musical traditions other than the Western classical one, and was common in the Western classical tradition itself until the early 19th century. Indeed, as Lydia Goehr (1992) argues, the whole idea of an integral musical work that is expected to be faithfully rendered by the performer is of recent vintage in the Western classical tradition. But to the extent that improvisation is *expected* and valued for itself, it is just not the case that the intention of the performer is—or should be—merely to carry out the intentions of the creator.

¹⁶ See Dipert (1986, 406 and 1993, 32-7) on the notion of a "virtual" individual agent.

- An antecedent, mental design phase is distinguished from a subsequent construction phase;
- The mental design is formed by an individual deliberating about the means for constructing a given product;
- The design specifies the features of the product and gives step-by-step instructions for constructing it;
- The construction phase is an unintelligent execution of these instructions carried out either by the designer or by some other individual or individuals.

These four features can be collectively characterized in a more general way in terms of a two feature model I shall call *centralized control*. The common sense of ‘control’ is that of a directive or determining force. But etymologically ‘control’ is an accounting term derived from the French ‘*contre-roller*,’ which refers to the practice of keeping a copy of a ledger for purposes of account verification. The etymological and common senses of ‘control’ are unified by the idea that in *faithful* copying the features of the original determine all the relevant features of the copy through the transcription process. And this is exactly how the accounts of production we have considered understand the relationship between mental design and product. Faithful copying involves two distinct phases—an antecedent phase in which an original is devised or identified, and a subsequent copying phase in which the features of the original are transcribed. These are the design and construction phases of the account of production we have been considering. The design (original) specifies all the relevant features of the product (copy).

In the case of simple copying procedures there is usually no need to specify exactly how to transcribe these relevant features—the copier can be assumed to know how to do this. But in cases where this assumption cannot be made, the original must also include instructions for how the copying is to be done. These instructions are the construction plan in the accounts of production we have been considering. The actual construction (copying), then, is a process that faithfully follows the instructions of the construction plan, and by so doing reproduces in a material medium the features of the product specified in the design. This faithful copying relationship between the design and construction phases of production is the *control* aspect of the model.

The main idea behind the *centralization* aspect of the model is that the design—and thus control—typically resides, or ideally should reside, in the mind of a single individual. This idea incorporates two related assumptions. First, even if the construction is carried out by a group of collaborating individuals, control of the construction process is implicitly understood to be exercised centrally by the designer. And second, the designer is implicitly understood to be a single individual. A common variation of this second assumption is to admit that design is often the work of a team or group of people, but to insist, as Dipert does, that it can be understood just as well or even better by reconstructing it as the work of a single individual. The centralization aspect of the model thus embodies the assumption that the paradigm case of production is the individual, skilled artisan working in solitary splendor.

This individualism with regard to the activity of production is the most important sense of centralization for our purposes here. But there is also a secondary sense of centralization found in many contexts. In the centralized control model the design is regarded as a mental plan

devised through deliberation. This implies centralization *within* the mind of the individual designer. Specifically, it implies a central planning function that takes as input all the available information about the desired product, processes it, and outputs a plan, or design.¹⁷

I said at the outset of this chapter that the two main problems with action theory from the point of view of its utility as the foundation for a theory of material culture are its lack of attention to collaboration on the one hand and to improvisation on the other. The historical review of Aristotle, Marx, and Dipert suggests that these *lacunae* are of long standing in the history of action theory, and that they are the consequences of an historical commitment to the centralized control model. My project in the rest of this chapter is to give some indication of how wide and deep these *lacunae* really are, and in consequence how inadequate the centralized control model really is.

It might seem strange that the existence of these *lacunae* has not been recognized, or when recognized has not been considered problematic for the centralized control model. But I think the implicit assumption has been that collaboration and improvisation are strictly secondary phenomena, and that once individual action and planning are well understood, accounts of collaborative action and improvisation will just fall out, or at worst require only a bit of tweaking and the reapplication of available theoretical resources. I fear this is unlikely to prove true.

First of all, the assumption that collaboration and joint action generally can be understood on the basis of a prior understanding of individual action and very little more, involves a commitment to some form of methodological individualism—the idea that social structures and group activities can be fully explained in terms of the intentional states and actions of the individuals making up the social group. Methodological individualism has been repeatedly challenged on the grounds that the individual is developmentally the product of social processes involving interactions with others, that individuals routinely carry out many of their activities in concert with others, and that even solitary action is continuously oriented to locally established practices and social norms. Thus individual action cannot be properly understood without a prior understanding of the relevant social practices and interactional structures.¹⁸ On some interpretations, this is taken to mean that it is individual action that is the secondary phenomenon and joint action that is primary. On the interpretation I favor, however, it means that joint action and individual action are equiprimordial modes of human activity which must be investigated side by side and in interrelation with each other from the start. In any case, if the critics of methodological individualism are correct, an adequate account of collaboration is not going to just fall out of a prior, adequate account of individual action. It will require a separate and parallel effort.

¹⁷ This assumption has been questioned recently in artificial intelligence and in cognitive science. On the one hand, there are problems with centralized planning systems in robotics, which tend to be brittle and unable to operate efficiently in real environments and real time. Systems that rely instead on a number of relatively independent modules each dedicated to producing a single, simple behavior do better in this respect. Although these decentralized systems often *look* like they have a central planner, this is in the eye of the observer and not in the actual operation of the system (Brooks 1999). On the other hand, there is evidence on a number of fronts that biological systems—real minds—are similarly decentralized. Cognition, in other words, is a distributed capacity and not a centralized one. As Andy Clark (1997, 33) puts it, the centralized, rational deliberator often turns out to be a distributed, adaptive responder instead.

¹⁸ For overviews of such critiques, see Joas (1996) and Kincaid (1986). More recently, and with more direct relevance to contemporary Anglo-American action theory, Margaret Gilbert (1989, 2000) has proposed what she calls ‘plural subject theory’ as an alternative to the individualistic analysis of group action.

Unfortunately, it is not clear that the centralized control model is on the way to an adequate account of individual action, either. The focus on control through planning as the sole structuring principle of human action betrays an implicit assumption that improvisation can be understood as a diluted or degraded form of planning, and that all the resources for understanding improvisation are thus in principle already available to the planning theorist. This assumption has not been as widely challenged as methodological individualism has been, although some challenges to it have surfaced in anthropology (Suchman 1987) and artificial intelligence (Agre and Chapman 1990, Brooks 1999), and have an ongoing influence in these fields. But this work is rarely cited in the philosophical literature,¹⁹ and no similar critique of planning has arisen indigenously in action theory. Indeed, improvisation is not even a topic of conversation among action theorists. But if the relevant work outside philosophy is on the right track, it indicates that the structure of improvised activity has distinctive features that set it apart from the structure of planned activity. Thus in this case, too, a separate and parallel effort is needed to understand improvisation and its relationship to planning.

In the next two sections I will lay the groundwork for investigating these issues by briefly describing some of the more distinctive general features of collaboration and improvisation. This basic phenomenology of collaboration and improvisation also provides an opportunity to introduce the two main example domains I will be focusing on throughout the book—cooking and song writing.

CENTRALIZATION AND COLLABORATION

Collaborative activity in the production of material culture is ubiquitous. Sometimes it is necessitated by the nature of the task and the agents involved. For example, collaboration is often necessary for constructing items of material culture because no single individual possesses the physical strength and/or the time to accomplish the construction task alone. Any item of material culture which is large (e.g., buildings, roads) or involves simultaneous performance (e.g., most plays and many musical pieces) requires collaborative construction for this reason. Many construction projects are also impossible to accomplish alone because no one individual possesses all the necessary knowledge or skills. Technically sophisticated items (e.g., computers, cars) are of this sort, but even in the case of buildings as relatively unsophisticated as residences, specialization to specific skilled tasks is standard. The carpenter, the plumber, the electrician, the mason, and the roofer are typically all different individuals. Collaboration among designers is also often necessitated by such natural limitations of individual knowledge and skill. This is perhaps most evident in the case of technically and socially sophisticated items of material culture like photocopiers or electric guitars. Companies that make such things usually employ design teams rather than individual designers for their products. Typically, some team members are responsible for the technical side, with further specialization for specific technical aspects; other team members are responsible for the aesthetic side and the marketing related aspects. An important variation on this last type of collaboration, and one that does not usually involve face-to-face contact between individuals, is the use of parts or materials that have been designed and

¹⁹ Clark (1997) is one of these rare cases. But Clark's work falls within the research area of cognitive science or philosophy of mind, not action theory. And in general, there seems to be comparatively little cross pollination between cognitive science and action theory.

constructed separately by other people. Cooks do not make their own butter or grind their own flour; and luthiers buy rather than make strings, metal fittings, and electric pickups for their instruments.

I want to stress the necessity of collaborative production under certain commonly encountered conditions in order to highlight the inevitability of collaboration in everyday activity. We can, and routinely do, entertain production projects we are incapable of carrying out alone; and we are social animals, so in principle we have collaborators available. Collaboration inevitably ensues. But it is important to note that for the most part collaboration is not strictly necessitated in this way. One reason for this is that even where it is possible for an individual to manage the physical requirements of a task alone or to learn all the necessary skills, there is often a clear advantage in relying on collaboration. Making your own butter is not all that difficult, but it takes time and equipment, so the most efficient thing is to get it from the dairy farmer instead. And then you have time and energy left over to make fancy pastries. This is precisely Marx's point when he says that the collaborating individual "strips off the fetters of his individuality and develops the capabilities of his species" (1976, 447). In addition, traditional divisions of labor and work specialization in human societies promote collaboration beyond what is either necessitated by the nature of the task and the limitations of individuals or recommended for reasons of efficiency. And last but not least, collaboration may be prompted not by hard and fast limits on what individuals can reasonably know or do by themselves, but by sheer differences of interest. For example, you might offer to make the salad for dinner because that is something you particularly enjoy doing, even though the main cook might be a fine salad maker and have plenty of time to do it.

Some important and especially interesting forms of collaboration arise because the designer and the constructor are often different individuals. Here, too, collaboration is frequently necessary. You would not expect a symphony composer or an architect to have either the physical capacity or all the requisite skills to realize their designs on their own. On the other hand, collaboration between designer and constructor is often in principle voluntary. Songwriters are typically quite capable of performing their own songs as solo pieces; and it would be an odd recipe deviser who could not prepare the dish themselves. But in both of these cases the proliferation of the item through collaboration between designers and constructors is integral to the social practices involved in the production and use of the items of material culture in question.

Collaboration between designers and constructors is especially significant with regard to the centralized control model, because it understands design and construction as distinct phases of production, with all the thinking relegated to the design phase and construction construed in consequence as merely unintelligent execution. This puts a particular slant on the nature of the collaborative relationship between designers and constructors—constructors collaborate by doing what designers tell them to do. But the discussion of Dipert's theory of performance already pointed to a problem with this slant. Although Dipert adheres to the idea of unintelligent execution as an ideal, he rightly stresses that it is an ideal rarely if ever realized in practice because the performer is not in a position to know the intentions of the designer in complete detail. And this means that the performer must be an intelligent executor to one extent or another. This point is not specific to performance theory, but has general application to any case where the collaborative relationship bridges design and construction. This will come up again in

the discussion of the control aspect of the centralized control model. But for now it should be noted primarily because it shows that the constructor, as a more or less intelligent executor, is a full fledged collaborator in the sense that some of the deliberation and creative decision making with regard to the features of the item of material culture are actually vested in the agents responsible for construction.

In addition, collaboration between designers and constructors is important because it calls attention to the fact that collaboration can be a more or less distant relationship. The designer and the constructor need not work face-to-face, know each other as individuals, or even communicate directly. Indeed, the designer may be long dead, as is often the case with the performance of musical pieces or plays. This kind of distant relationship between collaborators occurs in a slightly different form in collaboration within both the design phase and the construction phase. As we noted above, component parts of items of material culture are often separately constructed and simply incorporated into the item presently under construction. These component parts often are also separately designed, and either specified as integral components by the designer, or decided upon by the constructor. For example, a recipe writer might include sherry in the list of ingredients, but the recipe for the sherry and the making of it are entirely the work of other people.

A particularly interesting variant I shall call creative appropriation combines distant collaboration between designers and constructors with execution that is not merely intelligent but overtly creative. A common example is writing new lyrics to an existing tune—the Beach Boys’ “Surfin’ USA,” for instance, is Chuck Berry’s “Sweet Little Sixteen” with new lyrics written (probably) by Brian Wilson. More interesting cases involve the interpretive appropriation of a song in such a way that the music as well is substantially redesigned. Led Zeppelin did this brilliantly with Memphis Minnie’s “When the Levee Breaks,” for example, changing the tempo and the key, abandoning most of the original verses, rewriting some lines of the verses they kept, changing the melody substantially, and writing a bridge verse with a different melody and lyrics.²⁰

Creative appropriation may very well be more common in material culture domains that involve performances or items that are consumed during use, such as food and drink or medication. Such items of material culture have to be constructed anew for every occasion of use, and this means there is a permanent possibility of constructing them differently each time—changing the recipe a bit, or varying some aspect of the performance, for instance. But even more durable items of material culture require maintenance, repair and refurbishing, and may be voluntarily rebuilt or remodeled. Thus they too are subject to creative appropriation—barns or churches are sometimes remodeled for use as residences, for example.

At this point we must refine the distinction we made between the social and the collaborative when we examined Marx’s theory of production. Some of the dimensions of distance that may


²⁰ Chuck Berry’s “Sweet Little Sixteen” originally appeared on *One Dozen Berries* (Chess 1958). The Beach Boys’ retitled version is on the eponymous *Surfin’ USA* (Capitol Records 1963). Memphis Minnie’s original version of “When the Levee Breaks” can be found on *Kansas Joe and Memphis Minnie, 1929-1934* (Document 1991). Led Zeppelin’s version is on their untitled fourth album (Atlantic 1971). Both the Beach Boys and Led Zeppelin credited the originators of the appropriated songs. This is a good idea, as it can save you a lot in lawyer’s fees.

exist between collaborators have been mentioned. Collaborators may be geographically separated, temporally separated, communicatively isolated, or socially isolated in the sense that they do not know each other personally as individuals. So it looks like there is a continuum between the social and the collaborative, with clear cases at both extremes, but no sharp dividing line anywhere in between. A clear case of collaboration would be members of an extended family spending the afternoon in the kitchen putting up tomatoes. Here there is face-to-face contact between individuals in a single spatiotemporal location with ample use of direct communication. A clear case of sociality would be Big Joe Williams creative appropriation of a traditional field holler as the basis for his now ubiquitous “Baby, Please Don’t Go.” It is not clear exactly which traditional holler inspired Williams—different sources make different speculations. What is clear is that the traditional variants of this tune stretch back into the mists of time, and no one really knows who originated any of them. So Williams was relying on a social resource—a traditional tune with a number of variants, not attached to any specific place, date, or individual.²¹

When we consider this last example more closely, though, we can see that it is best described as a case of *individual* action, but with a social dimension. As the anthropologist, Peter Reynolds puts it:

Although human beings can and often do work alone, even solitary labor is almost always a social activity because it is typically directed towards social ends, requires materials conveyed through social exchange, is typically one step of a larger, cooperative endeavor, and uses skills developed as a member of society. (Reynolds 1993, 412)

But of course collaborative activities have this inherent social dimension as well. So although collaboration is overtly social in that it involves interaction among individuals, all human action is covertly and pervasively social in that it depends on social practices, products, norms, and so on. Thus the social does not first enter the arena of human action in overt forms of collaboration, but preexists both individual and collaborative action as a sustaining medium for human activity in general. Thus the examples in the previous paragraph mark points on a continuum between individual action and collaborative action, both of which are social throughout. In other words, there is no sharp dividing line between individual action and collaborative action, because the relationships and interactions between people that might be thought to constitute collaboration exist in varying degrees, from face-to-face interactions with persons known to you, to more distant interactions with strangers, or even, through their works, with people long dead and perhaps unknown even to history as named individuals.



Collaboration as a distinct species of sociality has a special significance for a philosophy of material culture, though. First, to the extent that action theory inquires into the social at all, the domain of the investigation has been collaborative action. Second, and more importantly,

²¹ Big Joe Williams recorded “Baby, Please Don’t Go” several times. One of these versions is on his *Complete Works, Volume I, 1935-1941* (Document 1991). Creative appropriation appears to have been Williams’ forte. The other thing he is famous for is the nine-string guitar he devised himself and played on many of his recordings.

collaboration is the most concrete manifestation of human sociality—sociality at work in the world, so to speak. And as sociality at work in the world collaborative activity is the growing point of human sociality in general and specific cultures in particular. Even more specifically, it is the primary locus of both the reproduction of material culture traditions and innovation within them.

This claim might be rejected by some on the grounds that individuals acting alone are at least equally responsible for the reproduction of culture; and that individuals acting alone are even more responsible for cultural innovation than collaborators, since innovation is the province of the lone artistic genius or solitary inventor. We have already stressed the ubiquity of design phase collaboration. But a more comprehensive answer to this objection is provided by Peter Reynolds (1993). He argues that the perception of individual action as central or especially typical of human activity in production contexts is a culturally specific bias, promoted as much by the scientific community as by anyone else, but not supported by the anthropological evidence. Reynolds compared videotapes of archaeologists demonstrating the construction and use of stone tools, in which the tool user is virtually without exception represented as an individual working alone, with field videotapes of Australian Aborigines making and using stone tools for everyday purposes, in which most of the activity is collaborative. He concludes that:

In industrial societies, ‘tool use’ is defined as something that is done with the hands by an individual working alone...and in popular books in Western culture, such as Jean Auel’s novels about the Paleolithic...the heroine discovers fire, invents the bow and arrow, tames the horse, and so on. Although this view of human history is dismissed by anthropologists as simplistic, it is not as far from canonical scientific theory as many archaeologists would like to believe. Indeed, evolutionary notions of tool use are permeated with Eurocentric assumptions that cannot be reconciled with the activities of children in a preschool, much less survive cross-cultural comparison.... (Reynolds 1993, 410)

Moreover, solitary labor is the exception in the human species rather than the rule, and in all societies work is typically performed by face-to-face task groups of people cooperating to accomplish a common goal. Even in the manufacture of stone tools, which can in fact be done by a single individual working alone (as indeed has been proven by the archaeologist turned flint-knapper), people such as the Australian Aborigines who depend on this technology for their livelihood nonetheless produce stone tools as a collective enterprise, as for example when one person molds the knife handles, while another attaches them to the blades. (Reynolds 1993, 412-13)

Reynolds’ main argument focuses on the twin claims that what is distinctive about human tool use is the social organization of activity in technical contexts; and that this social organization centers on collaborative face-to-face task groups engaging in what he calls ‘heterotechnic cooperation.’ In support of the first point he analyzed cross-species and cross-cultural videotapes, including videotapes of children playing at making things. The most striking difference he observed was that non-human primates invariably make and use tools individually whereas human beings normally make and use them in collaboratively organized groups. Other differences do exist. For example, human products tend to have a complex structure consisting of separable but interrelated parts; and human skill in manufacturing and manipulating these tools outstrips that of other primates. However, neither of these differences is particularly evident in children, whose skills and products are often well within the competence of non-human primates. On the other hand, children do typically exhibit full-fledged social and collaborative behavior in their play, which aligns them clearly with adult humans rather than with non-human primates. Reynolds does not deny that non-human primates do engage in social and collaborative behavior

such as mutual grooming and group play. His point is rather that in their case sociality and collaboration are *not* integrated into their technical activity as is the case with humans.

From this point of view, hominids diverged from apes when the social relations of grooming or play began to organize technical skills. (Reynolds 1993, 423)

In short, whatever other differences may exist between tool use in humans and non-humans, they depend on this basic and distinctive human integration of the social and the technical.

Collaborative production and use of material culture is thus not just typical of human beings as a matter of fact, but has deep phylogenetic roots.

Reynolds also has something to say about the exact nature of the collaborative organization introduced by the integration of the social and the technical.

I call this process *heterotechnic* cooperation ('different crafts') to emphasize the complementarity of social roles. Heterotechnic cooperation may be contrasted with *symmetric* cooperation, in which all the participants do the same thing at the same time in order to facilitate a common goal.... Thus human technology is not just 'tool use,' and not just 'cooperative' tool use, but tool use combined with a social organization for heterotechnic cooperation. This heterotechnic aspect of human tool use, characterized by complementary technical roles among the participants, is manifested in all human societies by a distinct form of social organization that I call the *face-to-face task group*—a social structure defined by *the shared intention to transform matter and energy through the cooperative and complementary use of tools and tool-using skills by a group of people in face-to-face contact*. (Reynolds 1993, 412)

The complementarity of social and technical roles brings with it a number of other features, such as communication among task group members. Reynolds also holds, somewhat controversially, that this entire complex of features rooted in the basic structure of heterotechnic cooperation is not only distinctively human, but characteristic of all human societies, from the less technologically sophisticated to the industrial.

From an examination of even a few minutes of videotape, it is clear that stone tool-making [by Australian Aborigines] is a cooperative process, not an example of an isolated craftsman working alone, and that even in this so-called primitive craft *the basic principles of a manufacturing system are present*: task specialization, symbolic coordination, social cooperation, role complementarity, collective goals, the logical sequencing of operations, and the assembly of separately manufactured parts. (Reynolds 1993, 411-12)

The evidence Reynolds adduces and the position he outlines on the basis of it pose two significant difficulties for the individualism of the centralized control model. First, if Reynolds is right, individualist approaches miss precisely what is most characteristic and distinctive about the production and use of material culture by human beings. Moreover, this casts particular suspicion on the idea, explicitly endorsed by Dipert, that individualist reconstruction of what are in fact collaborative activities constitutes an appropriate methodology for the philosophical study of material culture. Such a methodology conceals collaborative structures of interaction as a matter of principle. Thus if Reynolds is right this methodology is at best simplistic and at worst ideological, in import if not in intent.

Second, Reynolds' analysis of the social structure of collaborative interactions in production contexts suggests that individualist accounts will not provide adequate resources for the understanding of collaboration, because the central phenomena involved in collaborative activity, such as task specialization, complementarity of roles, and symbolic communication, simply do not show up in solitary individual activity. Sociality in general *is* common to the individual and the collaborative case. For example, a solitary thinker and a group might both use the same language. The problem is that the individual and the group do not necessarily use language in the same way or to accomplish the same things—a conversation is structurally quite a different sort

of enterprise from a monologue. For example, group deliberation—say, by a group of designers deliberating about the features of a new product—exhibits phenomena not found in individual deliberation. It is possible for the members of the group to surprise each other, for instance, either by bringing in information to which no one else had prior access, or by presenting a novel line of reasoning. Similarly, when group members disagree with each other, some standard ways of resolving disagreements include a senior member of the group having final say or veto power, majority rule implemented by voting, or agreement to disagree, perhaps with some members of the group splitting off to form a new group with a divergent project. Such phenomena indicate dimensions of group deliberation for which there are no good analogues in individual deliberation.²²

Similarly, there are questions about where the resulting mental design is in the case of a group. Presumably there is no group mind in which it is lodged. And in large part because of phenomena specific to collaboration like task specialization and role complementarity, it is not plausible that it is always present in its entirety in each member of the group. For example, if two people are collaborating to produce a birthday cake with frosted decoration, the person who is responsible for the cake part need know nothing about the frosting design or how to produce it. This even extends to not knowing the name and/or age of the person whose birthday it is—facts the frosting designer would have to know in order to follow the common practice of spelling out birthday wishes on the cake. In other words, and in more contemporary parlance, the design of the cake as a whole might be distributed over the group of collaborators. And at present the nature and operation of distributed representations is not well understood in general. Individualistic accounts of production committed to the centralized control model routinely ignore these and many similar issues about group production, regardless of whether the individualism is viewed as a matter of principle or as a matter of expository convenience.

CONTROL AND IMPROVISATION

The control ideal of the faithful copy calls for all relevant features of the product to be specified in the design along with a set of instructions for construction. In other words, the design is ideally supposed to be an algorithm (effective procedure) for realizing the product. Not surprisingly, this ideal is not an accurate description of everyday processes of production. The most recent thinking characterizes designs as plans.²³ And plans are sometimes thought of as recipes.²⁴ So in this section I will use the example of cooking with a recipe to demonstrate and categorize the divergences of everyday production from what would be expected based on the control ideal. Although they can take an entirely mental form, in literate cultures recipes

²² Social structures and processes are often used as models for less accessible individual psychological structures and processes. An important early example from philosophy is Socrates' use in *The Republic* of social justice in the city as a model for psychological justice in the individual soul (368c-369a). But it is agreed on all hands that the way such models in philosophy or science work does not entail that all the features of the model must have analogues in the phenomenon modeled. Indeed, typically only a selected few, theoretically relevant features of the model are expected to have analogues. So although it is quite legitimate and possibly very instructive to use group deliberation as a model for individual deliberation, this does not mean that going in the other direction, an account of individual deliberation will capture all the interesting and significant features of group deliberation.

²³ This is clear from the discussion of Dipert's work, but see also Houkes, Vermaas, Dorst, and Vries (2002) for an explicit application of planning theory to design of artifacts.

²⁴ This is often explicit in the artificial intelligence planning literature. See Pollack (1991, 44 ff.) for an example.

typically exist in a written form to which the cook refers. The examples I will focus on involve such written recipes, which may be conveniently regarded as externalized versions of the originating cook's mental design. For the originating cook the written recipe may serve as an aid to memory. But its main function is to communicate the design from the designer (the originating cook) to potential constructors (other cooks).

At first blush, recipes do seem to conform to the control model, since they normally include a description of the dish and instructions for combining the listed ingredients in order to achieve the result specified. But a closer look reveals an important divergence from control as an ideal—recipes typically leave many details open. Some of these involve features of the product. For example, a recipe may suggest either sour cream or yoghurt as a thinner for cucumber soup, or a cake recipe may simply tell the cook to 'frost with a butter frosting.' Many recipes also list some ingredients as optional altogether—chopped nuts are often optional in cookie recipes, for example. Indeed, this is such a common practice that some cookbooks have a convention for indicating optional ingredients, such as enclosing them in parentheses. In addition, recipes frequently specify ingredients generically. A recipe might call for a cup of shortening, for instance, and then you have to decide between butter, margarine, vegetable oil, and so on.

The construction instructions also leave details open. Consider the following (highly recommended) cookie recipe.

Rolled Pecan Cookies

4 tablespoons powdered sugar	1 tablespoon ice water
7 oz. butter (scant cup)	2 cups flour
2 cups pecans (small pieces)	1/8 teaspoon salt
1 teaspoon vanilla	

Cream butter and sugar, add the rest. Roll with palms of hands into finger lengths. Bake 45 minutes (325°F). Roll in powdered sugar while warm, or shake in bag with 1/2 cup powdered sugar.²⁵

Notice, first, that the order in which to mix ingredients is not completely specified—you are on your own after creaming the butter and sugar. Even where recipes do specify this order more completely, it is largely conventional. In making cookies and cakes it rarely matters whether you first sift the dry ingredients, then mix the wet ingredients, or *vice versa*. The only essential thing is that the dry and the wet ingredients be mixed separately before combining them. The instructions thus constitute at best a partial order on the steps in the construction process. Second, some crucial steps are not specified. Even a novice cook will know to put these cookies on a baking sheet before baking them, but the recipe does not actually tell you to do this. In addition, there is an open question about whether to grease the baking sheet or not. More experienced cooks will realize this is unnecessary because these cookies have so much butter in them they could not stick to anything if they tried. Finally, the last instruction explicitly requires the cook to decide between two options for coating the cookies with powdered sugar, and if the bag option is chosen the cook must further decide between paper and plastic.

²⁵ *The Settlement Cookbook*, 28th edition, revised and enlarged (Milwaukee: The Settlement Cookbook Company, 1947), 477.

So recipes diverge from the control ideal by designedly (as it were) not specifying all relevant features of the product or all relevant construction steps. In some cases the cook is expected to supply the requisite details on the basis of habits she has acquired (e.g., putting the cookies on a baking sheet). Here we see the importance of Aristotle's insistence on the necessity of having the right habits in addition to doing the right thinking. In other cases the cook is implicitly expected to make a decision based on background knowledge she may be presumed to have (e.g., greasing the baking sheet). This already indicates that some of the thinking necessary for making the product takes place in the construction phase. But in addition, recipes often explicitly prompt cooks to make decisions which will affect what features the product will have, or how those features will be achieved (e.g., listing optional ingredients or offering alternative construction procedures). This confirms that cooking practices do not regard construction as unintelligent execution. Rather they assign to the construction phase some of the deliberation and decision making regarded under the control ideal as the prerogative of the design phase.

This observation is confirmed and extended by looking at how cooks actually use recipes. The expectation generated by the control model is that cooks faithfully follow recipes to the maximum extent possible. But in fact cooks typically use recipes as a basis for improvisations of various sorts.²⁶ Improvisation is normally a response to local conditions. Sometimes these are difficulties encountered in the construction process. When you do not have an ingredient called for by a recipe, you can often substitute something else—cocoa and butter can be used in place of baking chocolate, for instance. On the other hand, sometimes local conditions serendipitously make available resources you can exploit. A cook with a walnut tree in his backyard might substitute walnuts for pecans in the recipe above. A third type of condition involves the special needs or desires of the cook and/or her clientele. In the cookie recipe above a vegan cook would substitute a vegetable shortening for the butter.

Cooks sometimes arrive at a stable customization of a recipe after a period of improvisatory experimentation. For example, because my oven runs a little hot, I tried out a number of different combinations of oven temperatures and baking times for the cookie recipe above, and finally settled on 325°F for forty minutes, plus turning the cookie sheet at the twenty minute mark for more even browning. On the other hand, recipes represent a permanent possibility of doing something different on the next occasion, perhaps even just on a whim rather than because of some specific difficulty or opportunity. A cook might just try rolling these pecan cookies in colored sprinkles or shaved chocolate instead of powdered sugar. In any case, the important point here is that in addition to all the details implicitly or explicitly left open by the recipe itself, the regular practice of cooks is to alter even details that *are* expressly specified to suit their own situations, purposes and desires.

²⁶ My thinking on this issue owes a lot to Suchman (1987), Agre and Chapman (1990), and Chapman (1991) and Agre (1997). I also want to stress that here I am talking about relatively experienced cooks. Novice cooks do tend to follow recipes rather more slavishly. But novice cooks by definition do not yet know how to cook—in particular, they do not yet have the fund of habits and background knowledge the experienced cook has accumulated. Consequently, there is no particular reason to think that experienced cooks are merely doing *better* what novice cooks do, and some good reasons to think they are doing something *different*. Specifically, the accumulation of knowledge and skill makes it possible for an experienced cook to use recipes as resources for improvisation rather than as controlling devices. See Dreyfus and Dreyfus (1986) for more on the dangers of trying to understand skilled activities in terms of what novices do.

So recipes diverge from the control ideal in two ways. They explicitly or implicitly require the construction phase cook to do some of the actual design work, and in addition they are routinely used by cooks as a basis for further improvisational alterations to, or extensions of, the design rather than faithfully followed. This means the constructor is not regarded as an unintelligent executor for the most part; and even to the extent that the constructor is so regarded on the face of it, the regular practice of cooks is to execute intelligently by revising recipes as need or opportunity arises.²⁷ These two characteristics of recipes and their use have been recognized in the action theory literature as characteristics of plans in general. Plans normally involve only partial and incomplete specifications of the goal and how to achieve it; and they are often revised by agents either prior to or during execution. Michael Bratman writes:

Our purposive agency is typically embedded in multiple, interwoven quilts of *partial*, future-directed plans of action. We settle in advance on such plans of action, *fill them in*, *adjust them*, and follow through with them as time goes by. (Bratman 1999, 1; emphasis added)

Thus it appears that the divergence of everyday designs and plans from the control ideal is commonly acknowledged. The process of getting from design to product—or more generally, from plan to action—is not a matter of faithful copying.

The reasons for this divergence are not far to seek. The full-fledged control ideal embodies a number of clearly false assumptions about what kind of world we live in and/or what kind of agents we are. Some of these assumptions were close to explicit in early artificial intelligence work on planning, as Martha Pollack describes it.

The agent is given a goal, it computes a plan for achieving it, and then, at least in principle, it executes that plan. The environment is quiescent; the agent is the only force acting on it. So nothing of significance happens while the agent is forming its plan. And nothing happens while the agent is executing that plan, except what the agent itself causes to happen. (Pollack 1992, 45)

This quiescence of the environment is the basic assumption that makes detailed advance planning and unintelligent execution even intelligible as an option. Only in a static world could an agent in principle count on just executing plans devised completely in advance. There is also an implicit correlative assumption here that the world is homogenous across agents such that if one agent devises a plan another agent can in principle always execute it.

But our actual world is obviously neither static nor homogenous. As Pollack goes on to explain:

Real environments are dynamic. They are populated by multiple agents that can and do effect change. Because they are dynamic, real environments may change while an agent is reasoning about how to achieve some goal, and these changes may undermine the assumptions upon which the agent's reasoning is based....

Real environments may also change while an agent is executing a plan, and again they may change in ways that make the plan invalid.... As if this were not enough, real environments may change in ways that do not invalidate a current plan, but instead offer new possibilities for action. (Pollack 1992, 46)

There is one possibility Pollack does not mention which would make it possible to deal with such dynamic environments and still live up to the control ideal. If agents were omniscient then the dynamic nature of the real environment would not matter because all possible changes could be predicted and prepared for in advance in the planning phase. But this assumption also is clearly false. Some divine agents may be omniscient, but human agents like ourselves must always

²⁷ Although we have been talking about cases where the originating cook and the constructing cook are different individuals, it should be noted that all of this applies equally well in cases where these roles are filled by the same individual. Indeed, the originating cook may have even fewer reasons for making detailed advance specifications for his own purposes, and even fewer inhibitions about departing from those specifications himself even after they have been set out in writing for others to use.

reckon with limited knowledge and limited cognitive capacities. So a serious commitment to the control ideal would commit us *pari passu* to false assumptions about the real nature of the world we live in on the one hand, or false assumptions about the kind of agents we are on the other hand. And of course no one does make these false assumptions. Even planning theorists in early AI work on action only assumed static environments transiently and for heuristic purposes.

But the consequences of rejecting these false assumptions underlying the control ideal have not been fully recognized. The control ideal is not just unattainable in practice—it is a false ideal. Trying to live up to it would encourage us to operate in a way that would fail most of the time, given the actual nature of our environment and of our agency. In real environments it does not make sense for even knowledgeable, experienced agents to devise detailed advance plans for unintelligent execution. This would militate against both flexible accommodation of problematic contingencies and flexible exploitation of serendipitous opportunities. Real human agents are much better off in real, dynamic environments with partial plans they can fill in as they go along and adjust or abandon to suit their changing circumstances. In some cases they may even be better off without any advance plans at all. And this means we are well advised to develop our skill at improvising rather than concentrating solely on our skill at constructing increasingly detailed plans and executing them faithfully. In other words, ideally we should aim to be better improvisers, not better controllers.²⁸

But for the most part there has been no turn away from control as an ideal, or, more specifically, away from planning as the foundation for the theory of action. In AI there was a brief movement away from planning in the late eighties, represented by researchers such as Rodney Brooks (1999), Philip Agre (1997), and David Chapman (1991), and there were a number of different ideas afloat in this period about viable frameworks for a non-planning based theory of action. But although the leaders of this movement have been very influential, the planning approach has remained predominant in AI. And in philosophical action theory no deviation from the planning paradigm is visible at all, despite widespread acknowledgement that plans are necessarily partial and subject to revision. This acknowledgement at least ought to have generated an interest in *how* we go about filling in, adjusting, and revising our plans, if not an interest in improvised action more generally—or so you would think. But this has not happened, and the reasons for this lie within action theory itself. Specifically, the way the planning paradigm has developed in action theory has made it seem as though once planning is thoroughly understood there will be nothing further to do. Any actions that might seem to fall under the heading of improvisation will already have been explained, and there will therefore be nothing for a separate theory of improvisation to do. Chapter Two investigates the current understanding of planning in action theory in order to show how this perception has developed, and why it is incorrect. Then in Chapter Three we return to the issue of collaboration and investigate the understanding of multiple agent action in current action theory. As we shall see, this understanding, too, is off the

²⁸ Vinod Goel (1995) makes a similar point. In producing designs which are themselves items of material culture (e.g., blueprints) there is a prolonged sketching phase in which alternative possibilities are generated, explored, and refined. Cognitive virtue here resides in producing sketches which are coarse-grained, ambiguous, subject to multiple interpretations, and so on, because these characteristics keep the design open-ended and easily transformable, which is what you want at this stage. His point—that cognitive virtue does not always reside in precision, elaborate detail, univocal designation, and so on, and in particular that it does not do so in design contexts—dovetails nicely with the claim here that improvisational activity is more virtuous than planning in production contexts.

mark in important respects. Finally, in Chapter Four we do some of the foundational work requisite to a more adequate action theoretic understanding of improvisatory collaboration.